TF-GNN: Graph Neural Networks in TensorFlow

Oleksandr Ferludin⋆¶, Arno Eigenwillig⋆¶, Martin Blais⋆†, Dustin Zelle⋆†, Jan Pfeifer⋆¶,
Alvaro Sanchez-Gonzalez⋆‡, Wai Lok Sibon Li⋆‡, Sami Abu-El-Haija†, Peter Battaglia‡,
Neslihan Bulut†, Jonathan Halcrow†, Filipe Miguel Gonçalves de Almeida†, Pedro Gonnet†,
Liangze Jiang◊, Parth Kothari¶, Silvio Lattanzi†, André Linhares†, Brandon Mayer†,
Vahab Mirrokni†, John Palowitch†, Mihir Paradkar¶, Jennifer She‡, Anton Tsitsulin†, Kevin Villela‡,
Lisa Wang‡, David Wong‡, Bryan Perozzi⋆†
tensorflow-gnn@googlegroups.com
¶: Google Core ML, †: Google Research, ‡: DeepMind, ◊: EPFL (work done at Google Research)
⋆: TF-GNN Top contributors

ABSTRACT

TensorFlow-GNN (TF-GNN) is a scalable library for Graph Neural Networks in TensorFlow. It is designed from the bottom up to support the kinds of rich heterogeneous graph data that occurs in today's information ecosystems. In addition to enabling machine learning researchers and advanced developers, TF-GNN offers low-code solutions to empower the broader developer community in graph learning. Many production models at Google use TF-GNN and it has been recently released as an open source project. In this paper we describe the TF-GNN data model, its Keras message passing API, and relevant capabilities such as graph sampling and distributed training.

1 INTRODUCTION

Machine Learning (ML) techniques have applications across domains as varied as medicine, social networks, biochemistry, robotics, and more. The success of many ML models is driven by their ability to incorporate different modalities of data (e.g. vision, text, sound, timeseries and geometric), each with its own unique structural (ir)regularities. Traditionally, software frameworks for machine learning (e.g. TensorFlow [1], PyTorch [32]) have focused on modeling one modality at a time, such as vision [25, 26] or natural language [10, 30, 36]. However, the development of graph representation learning [9, 18] and subsequent industry interest [2, 3, 6, 16, 21, 29, 31, 33, 40, 43] has motivated the need for better software frameworks for learning with *graph*-structured data.

In this paper we introduce TF-GNN[1], a Python framework that extends TensorFlow [1] with Graph Neural Networks (GNNs) [9, 18]: models that leverage graph-structured data. TF-GNN is motivated and informed by years of applying graph representation learning to practical problems at Google. In particular, TF-GNN focuses on the representation of *heterogeneous* graph data and supports the explicit modeling of an arbitrary number of relationship (edge) types between an arbitrary number of entity (node) types. These relationships can be used in combination with other TensorFlow components, e.g., a TF-GNN model might connect representations from a language model to those of a vision model and fine-tune these features for a node classification task. Many teams at Google run TF-GNN models in production. We believe this to be a direct consequence of TF-GNN's multi-layered API which is designed for
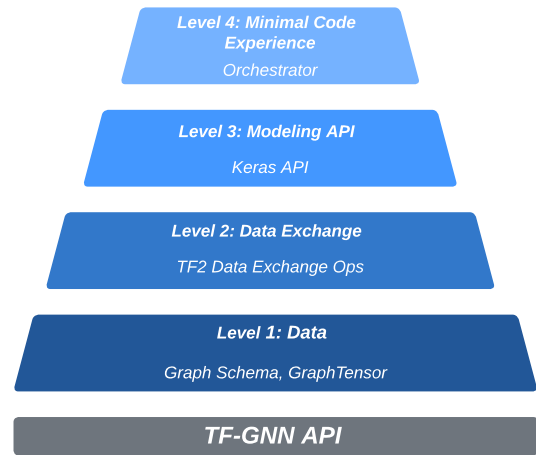
---

[1]https://github.com/tensorflow/gnn

Figure 1: TF-GNN's layered API decomposes the tools needed to create graph models into four distinct components of increasing abstraction.

accessibility to developers (regardless of their prior experience with machine learning).

Other software frameworks for learning from graph data have been proposed, most notably PyTorch Geometric (PyG) [12] and Deep Graph Library (DGL) [39]. We differ from DGL and PyG in three main ways. First, TF-GNN has been designed bottom-up for modeling heterogeneous graphs. Second, TF-GNN offers different levels of abstraction for increased modeling flexibility. Proficient users can leverage raw TensorFlow operations for message passing, limited only by their imagination. Intermediate users use the Keras modeling API and pre-built convolution layers, while beginner users can use the Orchestrator to quickly experiment with GNNs. Finally, TF-GNN is programmed on top of TensorFlow. As such, its goal is to support the many production-ready capabilities present in the TensorFlow ecosystem.

Interestingly, these differences can provide advantages for several use cases. For instance, (i) while it is possible for PyG to model heterogeneous graph data, its syntax advocates partitioning a heterogeneous graph into a set of homogeneous graphs. This makes it (programmatically) challenging to create a graph layer that pools from multiple node features at once, or even create new node or

feature types on the fly (i.e., through the network computation). However, TF-GNN's flexibility allows for aggregating from different node or edge types at once. Furthermore, (ii) TF-GNN offers *edge-centric*, *node-centric*, and *graph-centric* building blocks for GNNs. Existing frameworks (e.g., PyG) are node-centric, making implementing some models more tedious. On the other hand, edge-centric models, such as Graph Transformers [41] can be natively expressed in TF-GNN. Finally, (iii) TF-GNN's TensorFlow implementation inherits the benefits of the TensorFlow ecosystem, including access to model architectures for popular modalities (such as vision, text, and speech), and the ability to execute GNN models, for training and inference, on extremely fast hardware devices [20].

Summary of contributions. We present TF-GNN, an open-source Python library to create graph neural network models that can leverage heterogeneous relational data. TF-GNN enables training and inference of Graph Neural Networks (GNNs) on arbitrary graph-structured data. TF-GNN's four API levels allow developers of all skill levels access to powerful GNN models. Many TF-GNN models run in production at Google. Finally, as a native citizen of the TensorFlow ecosystem, TF-GNN shares its benefits, including pretrained models for various various modalities (e.g., a NLP model) and support for fast mathematical hardware such as Tensor Processing Units (TPUs).

2 OVERALL DESIGN

TF-GNN offers a layered API to build Graph Neural Networks that lets users trade off flexibility for abstraction. From least to most abstract, the layers (shown in Fig. 1) are:

- **API Level 1**: The *Data Level* (§3) takes care of representing heterogeneous graphs and loading them into TensorFlow, including technicalities like batching and padding.
- **API Level 2**: The *Data Exchange Level* (§4.1) provides operations for sending information across the graph between its *nodes*, *edges* and the *graph context*.
- **API Level 3**: The *Model Level* (§4.2) facilitates writing trainable transformations of the data exchanged across the graph in order to update the state of nodes, edges and/or context.
- **API Level 4**: The *Minimal-Code Experience Level* (§8.4) provides the Orchestrator, a toolkit for the easy composition of data input, feature processing, graph objectives, training, and validation.

This layered design is one reason for TF-GNN's successful adoption for graph models at Google. Users can start at a high level and only later go in deeper to tweak parts of the model. Some users may choose to only use the data level and its associated tooling (like the graph sampler, §6.1.1) and use their own modeling framework. The following sections describe the API levels in greater detail. Finally, we discuss other parts of the library designed for use in production models (§6).

3 TF-GNN HETEROGENEOUS DATA MODEL (API LEVEL 1)

To train a model on heterogeneous graph data, users first need to specify its node types, edge types and their respective features. This is done with the GraphSchema (§3.1). Based on that, the GraphTensor class (§3.2) can represent any graphs from the dataset.

3.1 Graph Schema

A GraphSchema object defines:

(1) One or more named **node sets** and their respective features.
(2) Zero or more named **edge sets** and their respective features. Each edge set has a specified source node set and a specified target node set. All edges in the set connect these node sets.
(3) **Context features**, which pertain to the entire input graph.[2]

As the name suggests, GraphSchema contains only an abstract definition of how entities are related (similar to an entity relationship diagram [27]) and has no actual data points. By definition, the node sets are disjoint, so they can serve as node types; same for edges. For each feature, the graph schema defines its name, its datatype (int, float, or string) and its shape, as in TensorFlow [1]. That means the graph as a whole is heterogeneous, but within each node or edge set, the features are uniformly typed and shaped. Each feature can have a different shape, so the features of a node set might comprise a scalar (say, a categorical feature), a variable-length sequence (e.g., tokenized text), a fixed-length vector (such as a precomputed embedding), a rank-3 tensor with an RGB image, and so on.

An example schema for a prototypical recommendation systems problem is shown in Figure 2a. It defines a heterogeneous graph structure with the following structure:

- Two node sets: "items" and "users". *Item* nodes have two features for their "category" (an integer scalar, corresponding to enumeration), and their "price" (a floating-point vector to hold item advertised prices). *User* nodes have three features, representing their "name" (string), "age" (int), and "country" (int).
- Two distinct sets of edges: "purchased" and "is-friend". *Purchased* edges connect users to items they have purchased, while *is-friend* edges connect users together.
- One context feature "scores", which applies to the graph as a whole.

3.2 GraphTensor

Figure 2b shows a graph that conforms to the example schema from above: Each circle corresponds to a node (colored by its node set), and the two types of lines correspond to the two different edge sets. Each node contains the features specified for its node set.

Figure 3 shows our approach to representing one such a graph in TensorFlow. We use zero-based indexing for the nodes in each node set and the edges in each edge set. Then any one feature on a node set or edge set can be represented by a tensor of shape $[n, f_1, \ldots, f_k]$, where $[f_1, \ldots, f_k]$, $k \geq 0$, is the feature shape from the GraphSchema. Moreover, for an edge set, the indices of source and target nodes are stored as two integer tensors of shape $[n]$ whose values are node indices in the node set specified by the graph schema.

The GraphTensor class expands on this approach to represent graphs as tensors through all stages of a TensorFlow program that builds a GNN model. Roughly, the stages are

(1) Reading, shuffling, batching and parsing GraphTensor values from tf.Example records on disk; possibly distributing them between replicas for data-parallel training.

---

[2]Section 3.2 will refine the notion of context for graphs merged from a batch of inputs.

(a) A graph schema.

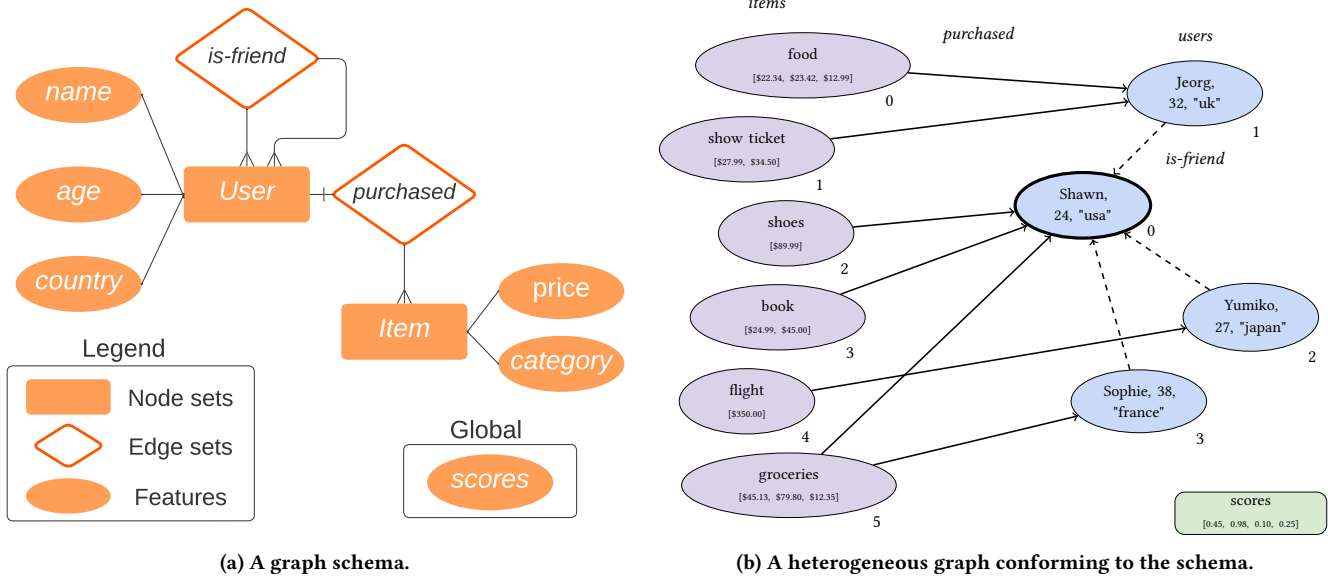(b) A heterogeneous graph conforming to the schema.

Figure 2: The recommending system example in the text uses the GraphSchema visualized in part (a). Part (b) shows a possible graph with the node sets, connecting edge sets, and features prescribed by the schema.
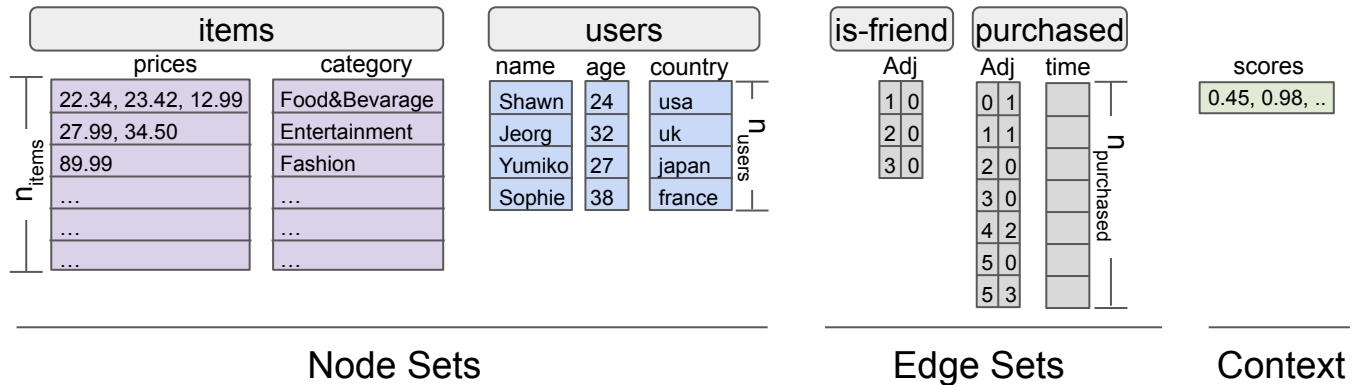

Figure 3: A `GraphTensor` to store the graph from Fig. 2b comprises tensors for all the features and for the adjacency data of each edge. (The size tensors are not shown here.) A Python expression for this `GraphTensor` object is shown in the Appendix.

(2) Transforming one or more input features per node or edge into a fixed-size representation for deep learning.

(3) Running a graph neural network for several rounds to update the hidden states of nodes (and possibly edges) from neighboring parts of the graph, followed by reading out the relevant hidden states and computing the model's output.

GraphTensor supports batching natively, and is indeed a tensor of graphs with a shape $[g_1, \ldots, g_r]$. A scalar GraphTensor has shape $[\,]$ and holds a single graph, while a GraphTensor of shape $[g_1]$ holds a vector of g_1 graphs, as usual for training with minibatches. Ranks $r > 1$ are rarely needed.

Each node set and edge set holds a dictionary of named features. Each feature is a tensor of shape $[g_1, \ldots, g_r, n, f_1, \ldots, f_k]$. GraphTensor allows that a feature dimension f_i may vary between the items of one node/edge set, or that the number n of items varies between

the multiple graphs in a GraphTensor. In both cases, the solution is to store the feature as a `tf.RaggedTensor`, not a `tf.Tensor`. Under the hood, this stores an explicit partitioning for each non-uniform, or "ragged", dimension. The shape reports its size as None.

To support feature processing, TF-GNN lets you create a new `GraphTensor` from an old one by replacing some or all of the features while keeping track of the implied schema change.

Finally, in service of training models, GraphTensor provides a method to merge a batch of inputs to a scalar GraphTensor. For each node/edge set, this method concatenates its elements across the batch of inputs and adjusts the node indices stored on edges correspondingly. The result is a GraphTensor of shape $[\,]$ with a flat index space $0, 1, \ldots, n_{\text{total}} - 1$ for each node/edge set, across the boundaries of batched examples. Features get the shape $[n_{\text{total}}, f_1, \ldots, f_k]$. Conveniently, such features can be represented as a `tf.Tensor`

if all feature dimensions are fixed, even if a node/edge set's size n_{total} is not constant between batches. If necessary, it can be made constant as well by adding a suitably sized padding graph to each batch of input graphs and assigning it weight 0 for training the GNN. Standard GNN operations on nodes and edges respect the boundaries between the merged input graphs, because there are no edges connecting them. To achieve the same for context features, GraphTensor supports the notion of **components** in a graph, and stores context features indexed by component. When a batch of graphs is merged, each input graph becomes one component of the result.

4 MODELING WITH TF-GNN

The core of TF-GNN is specifying how a computation utilizes graph structured data. In this section we detail the low level (TensorFlow) and high level (Keras) APIs used to construct GNNs.

4.1 Data Exchange Ops (API Level 2)

TF-GNN sends data across the graph as follows. *Broadcasting* from a node set to an edge set returns for each edge the value from the specified endpoint (say, its source node). *Pooling* from an edge set to a node set returns for each node the specified aggregation (sum, mean, max, etc.) of the value on edges that have the node as the specified endpoint (say, their target node.) The tensors involved are shaped like features of the respective node/edge set in the GraphTensor and can, but need not, be stored in it. Similarly, graph context values can be broadcast to or pooled from the nodes or edges of each graph component in a particular a node set or edge set.

Unlike multiplication with an adjacency matrix, this approach provides a natural place to insert per-edge computations with one or more values, such as computing attention weights [7, 37], integrating edge features into messages between nodes [13], or maintaining hidden states on edges [5].

4.2 Model Building API (API Level 3)

At API Level 3, TF-GNN follows standard TensorFlow practice and adopts Keras to express trainable transformations and their composition into models. API Levels 1 and 2 can serve other ways of modeling just as well. The shape $[n, \dots]$ of feature tensors allows to reuse standard neural network layers for item-wise transformations of node/edge sets, with set size n in place of a batch size.

A typical GNN model (cf. §3.2) consists of (i) feature transformations, (ii) a base GNN, and (iii) the final readout and prediction. TF-GNN lets you express this as a sequence of Keras layers that each take a GraphTensor input and return a GraphTensor output with transformed features – or a Tensor for reading out the final prediction.

4.2.1 Feature transformation layers. The feature transformations treat each node/edge set in isolation. Depending on the available features, they can range from simple numeric transformations to running an entire deep learning model to compute an embedding of, say, image or text data. TF-GNN lets you plug in other TensorFlow models and fine-tune them jointly while training the GNN on top. In the end, the representations of multiple input features are combined to form one initial `"hidden_state"` feature. TF-GNN's

MapFeatures layer makes it easy to build Keras sub-models for each node/edge set that map a features dict to a transformed features dict, and eventually the hidden state.

4.2.2 Graph Neural Network layers. The base GNN model is expressed as a sequence of *graph update* layers, each of which accepts a GraphTensor with `"hidden_state"` features and returns a new GraphTensor with these features updated. Each layer object has its own trainable weights; weight sharing can be achieved by using the same layer object repeatedly.

Users can define their own graph update layers, or reuse those from the growing collection of models bundled with the TF-GNN library (e.g., see the case study in Section 8). A user-defined graph update can, if needed, contain free-form code with an arbitrary composition of trainable transformations and broadcast/pool operations across all parts of the graph.

More commonly, graph updates are constructed from pieces operating on individual node sets or edge sets. TF-GNN provides a generic GraphUpdate class for that purpose, based on the following breakdown of updating a heterogeneous graph.

Consider any node v in some node set V of a heterogeneous graph. Starting from the initial hidden state $\mathbf{h}_v^{(0)}$, successive GraphUpdates compute the hidden state of v as

$$\mathbf{h}_v^{(i+1)} = \text{NextNodeState}_V^{(i+1)}(\mathbf{h}_v^{(i)}, \overline{\mathbf{m}}_{E_1,v}^{(i+1)}, \dots, \overline{\mathbf{m}}_{E_k,v}^{(i+1)}) \quad (1)$$

using the pooled messages $\overline{\mathbf{m}}_{E_j,v}^{(i+1)}$ received by node v in round $i+1$ along all edge sets $E_1, \dots, E_k$ incident to node set V. Notice that their number k is a constant from the graph schema for all $v \in V$.

Let $\mathcal{N}_{E_j}(v) = \{u \mid (u, v) \in E_j\}$ denote the neighbors of v along one edge set E_j, and notice that the size of $\mathcal{N}_{E_j}(v)$ may vary with $v \in V$. GraphUpdate supports two ways of computing $\overline{\mathbf{m}}_{E_j}$: in one step directly from the neighbor nodes as

$$\overline{\mathbf{m}}_{E_j,v}^{(i+1)} = \text{Conv}_{E_j}^{(i+1)}(\mathbf{h}_v^{(i)}, \{\mathbf{h}_u^{(i)} \mid u \in \mathcal{N}_{E_j}(v)\}), \quad (2)$$

or in two steps, materializing a per-edge message in the GraphTensor as

$$\begin{aligned} \mathbf{m}_{E_j,(u,v)}^{(i+1)} &= \text{NextEdgeState}_{E_j}^{(i+1)}(\mathbf{h}_u^{(i)}, \mathbf{h}_v^{(i)}, \mathbf{m}_{E_j,(u,v)}^{(i)}), \\ \overline{\mathbf{m}}_{E_j,v}^{(i+1)} &= \text{EdgePool}_{E_j}^{(i+1)}(\mathbf{h}_v^{(i)}, \{\mathbf{m}_{E_j,(u,v)}^{(i+1)} \mid u \in \mathcal{N}_{E_j}(v)\}). \end{aligned} \quad (3)$$

The two-step approach of Eq. (3) supports recurrence in $\mathbf{m}_{E_j,(u,v)}$, effectively turning it into a hidden state for edges. With that, and a context (or "global") state not shown here, this approach covers Graph Networks [5] and generalizes them to heterogeneous graphs.

Without recurrence in $\mathbf{m}_{E_j,(u,v)}$ (but possibly a constant edge feature in its place), this approach also covers Message Passing Neural Networks [13], generalized to heterogeneous graphs. The Conv abstraction from Eq. (2) is useful to express any of a number of successful GNN architectures in a single Python class and to transfer them directly to heterogeneous graphs with an arbitrary schema. Section 4.3 and the Appendix review some concrete cases, including Graph Convolutional Networks [24], which popularized the term *graph convolution* that we adopt here.

4.3 Implementing Popular Architectures

Graph Convolutional Networks [24]: GCNs for symmetric homogeneous graphs rely on adding loops (v, v) to the single edge set E to feed $\mathbf{h}_v^{(i)}$ into the computation of $\mathbf{h}_v^{(i+1)}$ along with the neighbor nodes. TF-GNN implements them by specializing Eq. (1) and (2) to

$$_v^{(i+1)} = {}_v^{(i+1)} = \sigma\Big(\sum_{u \in \mathcal{N}(v) \cup \{v\}} \frac{1}{\sqrt{d_u d_v}} \mathbf{W}^{(i+1)} \mathbf{h}_u^{(i)} \Big), \qquad (4)$$

where d_u is the in-degree of node u including loops and σ is an activation function, such as ReLU. Observe how Conv from Eq. (2) is the graph convolution, and NextNodeState trivially passes through its result.

The relational extension, **R-GCN**, [34] considers heterogeneous graphs and uses

$$\mathbf{h}_v^{(i+1)} = \sigma\Big(\sum_{j=1}^{k} \overline{\mathbf{m}}_{E_j, v}^{(i+1)} + \mathbf{W}_V^{(i+1)} \mathbf{h}_v^{(i)} \Big), \qquad (5)$$

$$\overline{\mathbf{m}}_{E_j, v}^{(i+1)} = \frac{1}{|\mathcal{N}_{E_j}(v)|} \sum_{u \in \mathcal{N}_{E_j}(v)} \mathbf{W}_{E_j}^{(i+1)} \mathbf{h}_u^{(i)}$$

with separate weights for each edge set and node set. This translates immediately to Conv and NextNodeState maps.

GraphSAGE [17] considers homogeneous sampled subgraphs and proposes several aggregator architectures, which translate directly to choices for Conv in Eq. (2). Its NextNodeState function turns out to be the special case $k = 1$ of Eq. (5) for R-GCN. In the GraphUpdate framework, GraphSAGE generalizes naturally to the heterogeneous case by running its Convs on multiple edge sets and combining them as in Eq. (5) for $k > 1$.

GAT [37] extends GCN by replacing the weighted sum in Eq. (4) by a concatenation of multiple weighted sums (attention heads), each with its own data-dependent weighting (formulas omitted for brevity). A Conv operation can naturally express GAT and other attention mechanisms on edges, including **GATv2** [7] and **Transformer-style dot-product attention** [7, 11, 23, 36].

The GraphUpdate framework allows to generalize GAT/GATv2 directly to the heterogeneous case, with no extra coding, analogous to the generalization from GCN to R-GCN. Attention is distributed separately between the edges of each edge set; learning the relative importance of different edge sets (relation types) is left to their separate weight matrices $\mathbf{W}_{E_j}^{(i+1)}$ in Eq. (5).

TF-GNN provides a base class for implementing Conv operations that allows a unified implementation of attention (i) from a node onto its neighbors, possibly combining the neighbor node state with a feature from the connecting edge; (ii) from a node onto its incoming edges; (iii) from the graph context onto all nodes; (iv) from the graph context onto all edges. Cases (ii–iv) provide attention for all aggregation steps of Graph Networks [5]. Appendix A.4 shows the unified implementation of GATv2 attention for all four cases.

5 ORCHESTRATION (API LEVEL 4)

At the highest level, TF-GNN provides the *Orchestrator*: a quick start toolkit with solutions for common graph learning tasks. It includes common graph learning objectives, distributed training
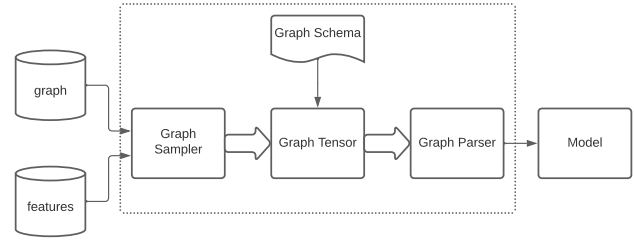
Figure 4: Diagram of massive-graph sampling and training pipeline with TF-GNN.

capabilities, accelerator support and the handling of many TensorFlow idiosyncrasies. The Orchestrator collects the tools necessary for (i) elevating the novice to a TF-GNN power user and (ii) increasing the scope of the graph learning expert's innovation. The toolkit supports graph learning research by offering both a standard framework for the reproduction of results and a shared catalog of SotA and convenience graph learning techniques and objectives. Specifically, the Orchestrator performs the:

(1) Reading input data to extract input graph(s) as `GraphTensor` instances $\mathbf{X}$ and corresponding label(s) $\mathbf{Y}$,
(2) Processing features for a specific dataset $\mathbf{X} \to \mathbf{X}'$,
(3) Adapting a model to the graph learning objective $\mathbf{M} \to \mathbf{M}'$,
(4) Training the adapted model $\mathbf{M}' \colon \mathbf{X}' \to \mathbf{H}$ to minimize the loss between $\mathbf{H}$ and $\mathbf{Y}$,
(5) Exporting the model $(\mathbf{M}')$ for inference or deployment.

The Orchestrator provides abstractions for the composition of these five steps.

- **Data source** (Python class `DatasetProvider`): An arbitrary source (on disk or in memory) that produces `GraphTensor` and a corresponding schema.
- **Feature processing** (`GraphTensorProcessorFn`): Feature manipulations for a `GraphTensor`–typically associated with a specific dataset.
- **Task** (Python class `Task`): Defines the graph learning objective on top of the base GNN model (e.g., node classification, or regression, or a self-supervised task like DeepGraphInfomax [38]) and its ancillary pieces like source data transformations.
- **Model**: An arbitrary model that operates on `GraphTensor`. The model is adapted to a graph learning objective by the above **Task**.
- **Training Hyperparameters**: Including the choice of optimization algorithm (e.g., Adam [4]), learning rate, as well as model hyperparameters (e.g., number of layers and their widths). Integrated with an automated hyperparameter tuning service [14, 35].

6 SAMPLING AND SCALING

At Google, we need to build neural network models for graphs of incredible size (trillions of edges). Heterogeneous graphs of this scale that have rich node and edge features cannot fit in the memory of single machines. TF-GNN offers a variety of approaches for transforming graph data. In our experience, the right approach typically depends on the graph's size. Here we discuss some details regarding processing datasets.

Algorithm 1 Distributed Sampling

Require: $S_\varnothing$: the initial set of nodes to sample
E: the edge sets to sample from
p: a sampling plan.
Ensure: $\mathcal{G}$, a set of Graph Tensors

$\text{frontier}_0 = \text{Sample}(S_\varnothing, E_{p_0})$
for $i = 1; i \leq p.\text{steps}$ **do**
  $\text{frontier}_i = \text{Sample}(\text{frontier}, E_{p_i})$
end for
$\text{edge\_groups} = \text{frontier.GroupBy(frontier.sample\_id)}$
$\text{edge\_groups} = \text{DeduplicateNodes(edge\_groups)}$
$\text{edges\_with\_features} = \text{lookup\_features(edge\_groups)}$
$\mathcal{G} = \text{create\_graph\_tensors(edges\_with\_features)}$

6.1 Sampling

On massive, well-connected graphs, even deep GNNs rarely aggregate features from distant nodes. As a result, when making a prediction on a root node (or a root activation node for edge/neighborhood prediction), it is frequently equivalent for the GNN to operate on a sufficiently-wide *subgraph* around the root [17]. For example, training and inference with a 2-layer GCN [24] only requires the 2-hop subgraph around the root.

TF-GNN formalizes this idea of rooted subgraph sampling with a *sampling plan*. The sampling plan describes the set of paths to sample from a graph schema in order to created rooted subgraphs. The key concepts of a sampling plan are: (1) the seeds – the nodes to start the expansion from, (2) the edge sets to expand out through, and (3) the kind of sampling to use in the expansion. This is discussed in detail during our case study in Section 8.

6.1.1 Large Scale: Distributed Sampling. One practical setting is the case of training GNNs on large scale data (e.g. graphs with *billions* of nodes). To train GNNs on massive graphs, we first run a distributed *subgraph sampling* operation which queries each root node for a subgraph, and writes each subgraph to disk as an individual GraphTensor.

The distributed sampling algorithm is presented in Algorithm 1. At a high level, it operates by repeatedly applying a distributed Sample function to grow the effective frontier of all node's samples at once. Each invocation of this function joins together a set of nodes to sample, and an edge set to sample over. After performing repeated sampling operations, the samples are joined by their sample id to form connected subgraphs. These subgraphs are then joined with the appropriate node features and converted to GraphTensors. Figure 4 shows an illustration of how sampling fits into the training pipeline.

The set of sampled GraphTensors is typically stored on distributed cloud storage. For the purpose of easing downstream processing, the GraphTensors are randomly grouped into file shards, which are then used as input for model training or inference.

6.1.2 Medium Scale: In-memory Sampling. Not all datasets are large enough to require distributed computation. In the case of smaller datasets (perhaps those with less than 100 million nodes),

the input graphs are small enough to fit into memory on one machine. However it can still be beneficial to perform sampling, in order to speed up training. In this case all graph nodes, edges, and features, are loaded for processing at once, and an in-memory sampling operation follows the sampling plan to generate GraphTensors on-the-fly. Here, one can express the objective function on the entire graph, e.g., as a node classification cross-entropy loss on all labeled nodes instead of piece-wise on rooted subgraphs. TF-GNN offers separate functionality and tutorials for these settings.

We note that unlike with distributed sampling, samples generated by the in-memory procedure are typically not persisted, and instead are used on-demand during the training or inference process.

6.1.3 Small Scale: No sampling. There are also many applications of TF-GNN where the graph(s) under consideration are quite small, and so each graph can just be completely contained inside a single GraphTensor. This setting is similar to the standard in-memory application of GNNs, and therefore we don't devote much space to it here.

6.2 Training

After creating GraphTensors, the next step is deciding how to train the model. TF-GNN's orchestration layer aims to make GNN training no different than normal model training. Specifically, it allows developers to provide a TensorFlow `tf.distribute.Strategy`[3] which allows both distributed training and accelerator (GPU, TPU) use.

6.2.1 Distributed Input Processing. TF-GNN's orchestration layer also integrates with TensorFlow's `tf.data` service[4], which allows utilizing a separate distributed compute cluster just for creating input for training, and performing CPU-based preprocessing. This can greatly accelerate TF-GNN model training by reducing I/O bottlenecks.

6.2.2 Model Serialization. After the training process is complete, a standard TensorFlow `SavedModel`[5] is created for the trained model.

6.3 Inference

Once we have a trained model, there are a variety of different options for performing inference. In the simplest case, we can do offline inference with batch processing using the SavedModel and a set of serialized GraphTensors (perhaps made through the same sampling process). More complex productionization settings can host the SavedModel on an inference service using TensorFlow Serving[6], or convert it to TFLite for inference on a mobile or edge device.[7] In this case, the GraphTensors used for input will have to be generated by the calling application (perhaps via the in-memory sampler).

---

[3]https://www.tensorflow.org/guide/distributed_training
[4]https://www.tensorflow.org/api_docs/python/tf/data/experimental/service
[5]https://www.tensorflow.org/api_docs/python/tf/saved_model
[6]https://www.tensorflow.org/tfx/tutorials/serving/rest_simple
[7]https://www.tensorflow.org/lite

7 RELATED WORK

Here we provide an overview of related work, which we separate into two broad categories: single-machine and distributed software frameworks. While it is more common for single-machine frameworks to be on the limits of research and innovation, industrial applications require distributed support and pose challenges rarely addressed by pure research applications.

Single-machine libraries. PyG (PyTorch-Geometric) [12] is a defacto standard framework for GNNs in the PyTorch ecosystem. PyG provides automatic batching support, GPU acceleration, and an interface to common graph learning datasets. Its performance is further enhanced by a set of optimized sparse GPU kernels tailored towards graph learning workloads. Spektral [15] is a prominent TensorFlow framework that follows Keras model building principles. It offers a similar experience to PyG in the TensorFlow ecosystem but without any batching support.

TF-GNN differs from these in many ways. For example, Spektral computes edge-centric functions (such as GAT) on minibatches by converting sparse adjacency matrices to dense ones. In general, these frameworks are not designed to scale to large graphs, but allow for easy experimentation by researchers on small graphs. On the other hand, TF-GNN's primary purpose is to scale to large graphs.

Distributed libraries. Deep Graph Library (DGL) [39] allows switching the backend platform between PyTorch, TensorFlow, and Apache MXnet with minimal code modifications. The DistDGL extension [42] enables efficient multi-machine training with DGL. Training is supposed to be done on a fleet of CPU-heavy instances connected in a cluster with a wide communication channel. DistDGL is mainly focused on reducing communication between the workers, each of which is performing sampling and training simultaneously. It partitions the input graph with METIS [22] and uses each partition as an example. Unlike DistDGL, TF-GNN's distributed training is more general. TF-GNN does not assume that the the data contains clusters, or that the graph structure fit into memory for partitioning via METIS.

Graph-Learn (formerly AliGraph) [44] is an open-source industrial graph learning framework built on top of TensorFlow. It is designed to natively handle large heterogeneous graphs, and it employs several techniques to facilitate large-scale training. Their distribution strategy relies on distributing the graph among worker machines, with a requirement that all worker machines must be alive at the same time: their training would stop if any worker machine fails. This differs from the distribution strategy of TF-GNN. In particular, TF-GNN samples a large graph into subgraphs using a resilient distributed system [8]. Similarly, TF-GNN can be used with the asynchronous distributed model training in TensorFlow, which is robust to machine failures.

Paddle Graph Learning (PGL) [28] is probably the most similar to TF-GNN. It is founded upon message passing over heterogeneous graphs. There are two notable differences between PGL and TF-GNN. First, PGL is more restrictive: each node must have a single feature (it is non-trivial to combine visual feature and textual feature, per node, per se) and dictates that all nodes must have the same feature dimensions. In contrast, TF-GNN support multiple features per node type (including ragged feature dimensions) and two node types can hold different features. Second, PGL uses Paddle as the computation backend whereas TF-GNN uses TensorFlow. From a practical sense, it is easier to find state-of-the-art (SotA) network architectures and pre-trained models in TensorFlow, which would make it easier to combine per-modality SotA models within a GNN.

8 EXAMPLE USE CASE: OGBN-MAG

Now we turn our attention to an end-to-end example illustrating how the components of TF-GNN interoperate within a 'real world' setting. As most applications of TF-GNN are proprietary, in this section we consider how a hypothetical developer at Google would use TF-GNN to solve a node classification problem using an open source dataset instead (OGBN-MAG). While the application here is admittedly contrived, the process is very illustrative of how TF-GNN is used internally.

8.1 Problem Identification

Our Googler begins their TF-GNN journey by identifying a product need. After consultation with their manager, team, and project stakeholders they determine that they desire to perform a node classification prediction for academic papers to determine the venue (journal or conference) in which each of the papers has been published. Having identified the machine learning task at hand, they collect labels for the training dataset and split it into three segments based on the time the papers were written ("train" (year<=2017), "validation" (year==2018), and "test"(year==2019)).

The next step in their TF-GNN journey is deciding on the appropriate graph structure to use for this classification problem. For this step, our Googler talks to data owners and determines what kinds of entities and relations are available for use. They refine this down to four kinds of nodes (papers, authors, institutions, and fields of study) and five kinds of edges to include in their model. [8] To close out this step, they encode these relationships in a Graph Schema (shown in Figure 5. The full protocol buffer definition of the schema is available in the Appendix A7.1.

This graph schema provides rich neighborhoods to author and paper, the two most numerous kinds of nodes, from which representations can be computed. Institutions and fields of study on the other hand have an unique 'id' feature that gets embedded and forms their node representation. This will allow the GNN model to train embedding tables for their representations over time, instead of computing a representation for them on the fly from their neighbors.

8.2 Input Generation

Next our Googler creates a sampling specification to encode how they would like to convert the graph for their model.

Since this is a node classification problem for papers, the user chooses to extract subgraphs, each subgraph centered around a paper node. For each seed paper node, the user wishes to sample other papers the seed paper cites, authors of cited papers, their

---

[8]We note that there are not too many node features included in this model, but want to make clear that in most real world settings there are an abundance of node (or edge) features possible.
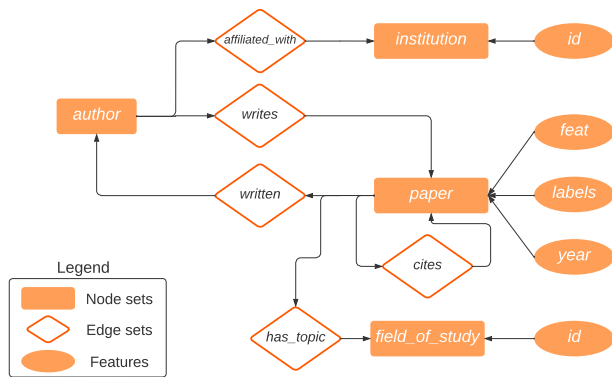
Figure 5: A Graph Schema chosen for papers classification on OGBN-MAG. It contains 4 node types and 5 edge types.

```python
# Each paper node is a seed for graph sampling.
seed_paper = SamplingSpecBuilder(schema,
                                 RANDOM_UNIFORM).seed(
                                 'paper')
# From each seed paper, sample cited papers.
cited_papers = seed_paper.sample(32, 'cites')
# From each paper (seed/cited), sample up to 8 authors.
authors = cited_papers.join([seed_paper]).sample(8,
                                                 'written')
# From these authors, sample up to 16 extra papers
    written by each.
author_papers = authors.sample(16, 'writes')
# From these authors, sample their affiliations.
_ = authors.sample(16, 'affiliated_with')
# From each paper (seed/cited/written), sample topics.
_ = author_papers.join([seed_paper,
                        cited_papers]).sample(16,
                        'has_topic')
# Create the SamplingSpec protocol buffer.
sampling_spec = seed_paper.build()
```

Figure 6: Sampling Specification for OGBN-MAG

affiliations, and other papers written by them. For all sampled papers (seed, cited, authored), user wants to also sample some fields of study. This sampling logic can be concisely written using TF-GNN's SamplingSpecBuilder – see Algorithm 6. The output of this sampling specification is encoded in a protocol buffer (Appendix A.7.2).

From the SamplingSpec that gets generated by Algorithm 6, our Googler user can choose which sampling pipeline to utilize for training a GNN. In this case, let us suspend disbelief momentarily and say the user decides that their graph data is too large for a single machine – as this is commonly the case for real datasets at Google. Therefore, they run the distributed sampler (described in Section 6) and generate a set of samples from the nodes in their graph.

8.3 Modeling

The next step in our Googler's journey is perhaps the most exciting part – the modeling phase! Here, the user must choose graph convolutions to transform data over their sampled input graphs.

After making some assumptions about the kind of information transfer necessary for the problem (and perhaps after performing a

consultation with the TF-GNN team), our user decides that a simple message passing neural network (MPNN) [13] is a promising first approach to their problem.

For exposition, let us start by writing an MPNN from scratch, before we look at the ready-to-use model pieces shipped with TF-GNN. The two key pieces are (i) a Conv layer, to define how messages are computed on one edge set and aggregated at its receiver nodes, and (ii) a NextState layer, to define how node states are updated based on the aggregated messages from the various incoming edge sets. Figure 7 shows the code to define a message passing neural network, with some boilerplate omitted. As you can see, there are 4 rounds of message passing, each implemented by a GraphUpdate layer that breaks it down into convolutions on the various edge sets and node state updates for those node sets that are not backed by embedding tables. This compact yet explicit representation of all trainable model pieces allows our user a great deal of control, separately for each edge set and node set.

The likely next step for our user is to add dropout, weight regularization and layer normalization to the MPNN shown above. This is straightforward to add to the Keras classes shown above, albeit verbose. This is where TF-GNN's bundled model collection comes in handy: The code in Figure 7 can be replaced by the much shorter Figure 8.

Beyond VanillaMPNN, TF-GNN offers models such as Graph-SAGE, GATv2 and (Transformer-style) MultiHeadAttention. These can be used wholesale, analogous to VanillaMPNN above, or mixed and matched per edge set / node set by combining them in one GraphUpdate.

After finishing a modeling implementation (either by writing one, or using a prebuilt one), our Googler is ready to set it up for training.

8.4 Orchestration

In order to add a loss to the model and set up training, our Googler then turns to writing a small configuration using TF-GNN's orchestration layer. This involves loading the data, choosing a Task for the model (RootNodeMulticlassClassification), and running the Keras trainer. Appendix A.6.4 shows a minimal but working TensorFlow training script that plugs the aforementioned code pieces into the `runner.run()` entry point.

Not shown here is the Runner's support for validation, checkpointing, padding inputs to fixed sizes (as required for Cloud TPUs), customizable export to a SavedModel for inference, model attribution utilities, unsupervised graph objectives and more. Please refer to the more comprehensive for detailed information.

8.5 Model Tuning

Finally, our Googler is ready to tune hyper-parameters to find the best model for their task. TF-GNN's support for Vizier [14, 35] hyper-parameter optimization allows them to explore different parameters including model parameters such as the dimensionality of the messages used and regularization parameters like dropout.

The Googler then conducts an extensive vizier study (here 100 trials where each trial is run for 20 epochs) with the objective is to maximize the accuracy on OGBN-MAG's validation set. Full details of the hyper-parameter search settings, along with the results, are

```python
class MyConv(tf.keras.layers.Layer):
  """Sends messages on an edge set and pools them for each
      receiver node."""
  def __init__(self, units, **kwargs):
    super().__init__(**kwargs)
    self._message_fn = tf.keras.layers.Dense(units,
        activation="relu")


  def call(self, graph, *, edge_set_name):
    # NOTE: The receiver is the source node from which the
        edge was sampled.
    receiver_states = tfgnn.broadcast_node_to_edges(
        graph, edge_set_name, tfgnn.SOURCE,
        feature_name=tfgnn.HIDDEN_STATE)
    sender_states = tfgnn.broadcast_node_to_edges(
        graph, edge_set_name, tfgnn.TARGET,
        feature_name=tfgnn.HIDDEN_STATE)
    message_inputs = tf.concat([sender_states,
        receiver_states], axis=-1)
    messages = self._message_fn(message_inputs)
    pooled_messages = tfgnn.pool_edges_to_node(
        graph, edge_set_name, tfgnn.SOURCE, "sum",
            feature_value=messages)
    return pooled_messages


class MyNextNodeState(tf.keras.layers.Layer):
  """Computes next node state from pooled messages and
      previous node state."""
  def __init__(self, units, **kwargs):
    super().__init__(**kwargs)
    self._update_fn = tf.keras.layers.Dense(units, "relu")


  def call(self, inputs):
    prev_node_state, edge_set_inputs, context_input = \
        inputs
    assert not context_input, "Unused_in_this_example."
    combined_inputs = tf.concat([
        prev_node_state
    ] + [
        pooled_messages
        for edge_set_name, pooled_messages in
            edge_set_inputs.items()
    ], axis=-1)
    next_node_state = self._update_fn(combined_inputs)
    return next_node_state


def model_fn(graph_tensor_spec):
  """Creates the GNN model from an input spec."""
  hidden_state_dim = 128
  message_dim = 128
  input_graph_tensor = tf.keras.layers.Input(type_spec=
      graph_tensor_spec)
  graph_tensor = tfgnn.keras.layers.MapFeatures(
      node_sets_fn=set_initial_node_states)(
          input_graph_tensor)
  for _ in range(4):
    graph_tensor = tfgnn.keras.layers.GraphUpdate(
        node_sets={
        "paper": tfgnn.keras.layers.NodeSetUpdate(
            {"cites": MyConv(message_dim),
             "written": MyConv(message_dim),
             "has_topic": MyConv(message_dim)},
            MyNextNodeState(hidden_state_dim)),
        "author": tfgnn.keras.layers.NodeSetUpdate(
            {"writes": MyConv(message_dim),
             "affiliated_with": MyConv(message_dim)},
            MyNextNodeState(hidden_state_dim)),
        })(graph_tensor)
  return tf.keras.Model(input_graph_tensor, graph_tensor)
```

Figure 7: MPNN implementation via GraphUpdate

```python
from tensorflow_gnn.models import vanilla_mpnn

def model_fn(graph_tensor_spec):
  ...  # As above.
  for _ in range(4):
    graph_tensor = vanilla_mpnn.VanillaMPNNGraphUpdate(
        units=..., message_dim=..., receiver_tag=tfgnn.
            SOURCE,
        l2_regularization=..., dropout_rate=...,
        use_layer_normalization=True,
    )(graph_tensor)
  return tf.keras.Model(input_graph_tensor, graph_tensor)
```

Figure 8: MPNN reuse from model library

	# params	validation	test
HGT (leaderboard)	26.8M	0.5124 ± 0.0046	0.4982 ± 0.0013
MPNN (tf-gnn)	**5.89M**	**0.5149 ± 0.0019**	**0.5027 ± 0.0022**

Table 1: The 'simple' model found here through hyper-parameter optimization outperforms a much higher capacity model from the OGB leaderboards.

found in Appendix A.6.3. Happily, they find a well-performing model with performance characteristics that meet their product requirements.

Table 1 shows how the the results from the Googler's Vizier study compare to a model with much more capacity (the Heterogeneous Graph Transformer [19]) from the OGBN-MAG Leaderboard. Interestingly, we see that a simple model (with no customization) and some hyper-parameter optimization can out compete a much more complicated model utilizing transformer-style attention. The TF-GNN platform aims to minimize the hurdles to placing these (and much more complicated) graph models into production settings.

9 CONCLUSIONS

In this work we have presented TF-GNN, our open source framework used for Graph Neural Networks at Google. TF-GNN is a software framework which reduces the technical burden for GNN productization and facilitates experimentation with GNNs. TF-GNN's expressive modeling capability allows complex relationships between nodes, edges, and graph-level elements in a model. This enables the straightforward implementation of intricate models. TF-GNN offers four levels of increasingly abstract APIs, serving a range of use-cases from beginners with a graph problem but little ML experience to ML researchers who desire complete control of their graph learning system. Many models at Google already use TF-GNN, and we believe that this project will accelerate the industrial adaptation of these promising models at more organizations.

REFERENCES

[1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. TensorFlow: A System for Large-Scale Machine Learning. In *12th USENIX symposium on operating systems design and implementation (OSDI 16)*. 265–283.

[2] Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. 2019. Mixhop: Higher-order graph convolutional architectures via sparsified neighborhood mixing. In *international conference on machine learning*. PMLR, 21–29.

[3] Rami Al-Rfou, Bryan Perozzi, and Dustin Zelle. 2019. Ddgk: Learning graph representations for deep divergence graph kernels. In *The World Wide Web Conference*. 37–48.

[4] Jimmy Ba and Diederik Kingma. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*.

[5] Peter Battaglia, Jessica Blake Chandler Hamrick, Victor Bapst, Alvaro Sanchez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, Caglar Gulcehre, Francis Song, Andy Ballard, Justin Gilmer, George E. Dahl, Ashish Vaswani, Kelsey Allen, Charles Nash, Victoria Jayne Langston, Chris Dyer, Nicolas Heess, Daan Wierstra, Pushmeet Kohli, Matt Botvinick, Oriol Vinyals, Yujia Li, and Razvan Pascanu. 2018. Relational inductive biases, deep learning, and graph networks. *arXiv/1806.01261* (2018).

[6] Aleksandar Bojchevski, Johannes Klicpera, Bryan Perozzi, Amol Kapoor, Martin Blais, Benedek Rózemberczki, Michal Lukasik, and Stephan Günnemann. 2020. Scaling graph neural networks with approximate pagerank. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2464–2473.

[7] Shaked Brody, Uri Alon, and Eran Yahav. 2022. How Attentive are Graph Attention Networks?. In *International Conference on Learning Representations*.

[8] Craig Chambers, Ashish Raniwala, Frances Perry, Stephen Adams, Robert R Henry, Robert Bradshaw, and Nathan Weizenbaum. 2010. FlumeJava: easy, efficient data-parallel pipelines. *ACM Sigplan Notices* 45, 6 (2010), 363–375.

[9] Ines Chami, Sami Abu-El-Haija, Bryan Perozzi, Christopher Re, and Kevin Murphy. 2022. Machine Learning on Graphs: A Model and Comprehensive Taxonomy. *Journal of Machine Learning Research* 23, 89 (2022), 1–64.

[10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 4171–4186.

[11] Vijay Prakash Dwivedi and Xavier Bresson. 2020. A Generalization of Transformer Networks to Graphs. *arXiv preprint arXiv:2012.09699* (2020).

[12] Matthias Fey and Jan E. Lenssen. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.

[13] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. 2017. Neural Message Passing for Quantum Chemistry. In *International Conference on Machine Learning*.

[14] Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Karro, and David Sculley. 2017. Google Vizier: A service for black-box optimization. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*. 1487–1495.

[15] Daniele Grattarola and Cesare Alippi. 2021. Graph neural networks in tensorflow and keras with spektral (application notes). *IEEE Computational Intelligence Magazine* 16, 1 (2021), 99–106.

[16] Jonathan Halcrow, Alexandru Mosoi, Sam Ruth, and Bryan Perozzi. 2020. Grale: Designing Networks for Graph Learning. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (Virtual Event, CA, USA) *(KDD '20)*. Association for Computing Machinery, New York, NY, USA, 2523–2532. https://doi.org/10.1145/3394486.3403302

[17] W. Hamilton, R. Ying, and J. Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *Advances in Neural Information Processing Systems*.

[18] William L. Hamilton. 2020. *Graph Representation Learning*. Morgan & Claypool Publishers.

[19] Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. 2020. Heterogeneous Graph Transformer. In *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*. ACM / IW3C2, 2704–2710. https://doi.org/10.1145/3366423.3380027

[20] Norman P Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, et al. 2017. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th annual international symposium on computer architecture*.

[21] Amol Kapoor, Xue Ben, Luyang Liu, Bryan Perozzi, Matt Barnes, Martin Blais, and Shawn O'Banion. 2020. Examining covid-19 forecasting using spatio-temporal graph neural networks. *arXiv preprint arXiv:2007.03113* (2020).

[22] George Karypis and Vipin Kumar. 1998. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on scientific Computing* 20, 1 (1998), 359–392.

[23] Dongkwan Kim and Alice Oh. 2022. How to Find Your Friendly Neighborhood: Graph Attention Design with Self-Supervision. *arXiv preprint arXiv:2204.04879* (2022).

[24] Thomas Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*.

[25] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Proc. of the Advances in Neural Information Processing Systems (NIPS)*.

[26] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. 1998. Gradient-based learning applied to document recognition. *Proc. of the IEEE* 86, 11 (1998), 2278–2324.

[27] Qing Li and Yu-Liu Chen. 2009. Entity-relationship diagram. In *Modeling and Analysis of Enterprise and Information Systems*. Springer, 125–139.

[28] Yanjun Ma, Dianhai Yu, Tian Wu, and Haifeng Wang. 2019. PaddlePaddle: An open-source deep learning platform from industrial practice. *Frontiers of Data and Computing* 1, 1 (2019), 105–115.

[29] Elan Sopher Markowitz, Keshav Balasubramanian, Mehrnoosh Mirtaheri, Sami Abu-El-Haija, Bryan Perozzi, Greg Ver Steeg, and Aram Galstyan. 2021. Graph Traversal with Tensor Functionals: A Meta-Algorithm for Scalable Learning. In *International Conference on Learning Representations*. https://openreview.net/forum?id=6DOZ8XNNfGN

[30] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*.

[31] John Palowitch, Anton Tsitsulin, Brandon Mayer, and Bryan Perozzi. 2022. GraphWorld: Fake Graphs Bring Real Insights for GNNs. *arXiv preprint arXiv:2203.00112* (2022).

[32] Adam Paszke, Sam Gross, Francisco Massa, and Others. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*.

[33] Benedek Rozemberczki, Peter Englert, Amol Kapoor, Martin Blais, and Bryan Perozzi. 2021. Pathfinder discovery networks for neural message passing. In *Proceedings of the Web Conference 2021*. 2547–2558.

[34] Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2018. Modeling Relational Data with Graph Convolutional Networks. In *Proc. ESWC (LNCS, Vol. 10843)*, Aldo Gangemi et al. (Eds.). Springer, 593–607.

[35] Xingyou Song, Sagi Perel, Chansoo Lee, Greg Kochanski, and Daniel Golovin. 2022. Open Source Vizier: Distributed Infrastructure and API for Reliable and Flexible Blackbox Optimization. In *Proc. AutoML 2022 (PMLR, Vol. 188)*, Isabelle Guyon et al. (Eds.).

[36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*.

[37] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.

[38] Petar Veličković, William Fedus, William L. Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. 2019. Deep Graph Infomax. In *International Conference on Learning Representations*.

[39] Minjie Wang, Da Zheng, Zihao Ye, Quan Gan, Mufei Li, Xiang Song, Jinjing Zhou, Chao Ma, Lingfan Yu, Yu Gai, Tianjun Xiao, Tong He, George Karypis, Jinyang Li, and Zheng Zhang. 2019. Deep Graph Library: A Graph-Centric, Highly-Performant Package for Graph Neural Networks. *arXiv preprint arXiv:1909.01315* (2019).

[40] Minji Yoon, John Palowitch, Dustin Zelle, Ziniu Hu, Russ Salakhutdinov, and Bryan Perozzi. 2022. Zero-shot Transfer Learning within a Heterogeneous Graph via Knowledge Transfer Networks. In *Advances in Neural Information Processing Systems*, Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (Eds.).

[41] Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J Kim. 2019. Graph Transformer Networks. In *Advances in Neural Information Processing Systems*.

[42] Da Zheng, Chao Ma, Minjie Wang, Jinjing Zhou, Qidong Su, Xiang Song, Quan Gan, Zheng Zhang, and George Karypis. 2020. DistDGL: distributed graph neural network training for billion-scale graphs. In *2020 IEEE/ACM 10th Workshop on Irregular Applications: Architectures and Algorithms (IA3)*. IEEE, 36–44.

[43] Qi Zhu, Natalia Ponomareva, Jiawei Han, and Bryan Perozzi. 2021. Shift-Robust GNNs: Overcoming the Limitations of Localized Graph Training data. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 27965–27977. https://proceedings.neurips.cc/paper/2021/file/eb55e369affa90f77dd7dc9e2cd33b16-Paper.pdf

[44] Rong Zhu, Kun Zhao, Hongxia Yang, Wei Lin, Chang Zhou, Baole Ai, Yong Li, and Jingren Zhou. 2019. AliGraph: a comprehensive graph neural network platform. *Proceedings of the VLDB Endowment* 12, 12 (2019), 2094–2105.

A   APPENDIX

A.1   Graph Tensor Example

Consider again the example schema and graph provided in Figures 2 and 3. This example graph would be represented with the following tensors:

- Node feature tensors:
  - "items" category, string vector tensor: `["food", "show ticket", "shoes", "book", "flight", "groceries"]`, with shape `[6]`
  - "item" price, float tensor: `[[22.34, 23.42, 12.99], [27.99, 34.50], [89.99], [24.99, 45.00], [350.00], [45.13, 79.80, 12.35]]`, with shape `[6, None]`; an instance of `tf.RaggedTensor`
  - "users" name, string vector tensor: `["Shawn", "Jeorg", "Yumiko", "Sophie"]`, with shape `[4]`
  - "users" age, int vector tensor: `[24, 32, 27, 38]`, with shape `[4]`.
  - "users" country: int vector tensor: `[3, 2, 1, 0]`, with shape `[4]`, assuming that the country vocabulary enumeration is `dict(france=0, japan=1, uk=2, usa=3)`.
- Edge feature tensors:
  - "purchased" with
    * source indices, int vector tensor: `[0, 1, 2, 3, 4, 5, 5]`.
    * target indices, int vector tensor: `[1, 1, 0, 0, 2, 3, 0]`, both with shape `[7]`.
  - "is-friend" with
    * source indices, int vector tensor: `[1, 2, 3]`.
    * target indices, int vector tensor: `[0, 0, 0]`, both with shape `[3]`.
- Graph-level feature tensors:
  - "scores", float matrix tensor: `[[0.45, 0.98, 0.10, 0.25]]` with shape `[1, 4]`.

Note how the edges connectivity is encoded as source and target indices in the arrays of node features. The indices are indicating, for each edge, which position in the node feature array they are referring to. For example, the fifth values of the 'purchased/#source' and 'purchased/#target' is '[4, 2]', which link together nodes '"flight"' and '"Yumiko"'.

A.1.1   GraphTensor Features & Shapes. Shapes of feature tensors, stored in `GraphTensor`, are crucial for building models in TF-GNN. Each tensor has a particular shape constraint based on its containing node set, edge set, or context dictionary. The shapes are as follows:

- Node features. All the features associated with a node set share the initial dimension, which is the total number of nodes in the node set. In the example above, features for node "items" have 6 nodes, and so all their tensor shapes are of the form `[6, ...]`. In general, node features will have the `[num_nodes, feature...]` shape. For example, a simple 64-dimensional embedding of each node results in shape `[6, 64]`, while a 224x224 image with 3 color channels stored at each node could result in `[6, 224, 224, 3]`.
- Edge features and indices. Similarly, all the features associated with an edge set shares the leading dimension, which is the total number of edges in the edge set. This includes the edge indices (rendered above as special features `source` and `target`). In the example, features for edges "purchased", both of these have shape `[7]`. If edges have information encoded as features, *e.g.*, an embedding of shape `[32]`, then the edge feature tensor would be of shape `[7, 32]`. In general, edge features will have shape `[num_edges, feature...]`.
- Context features. Context features apply to one component of the graph. An input graph parsed with a GraphSchema has a single component, so its context features have shapes `[1, ...]`. After batching inputs and merging batches of inputs into components of a single graph (see §3.2), there are multiple components, and their number appears as the outermost dimension in the shape of context features.

A.2   Graph Tensor API

A.2.1   GraphTensor Interface. The `GraphTensor` objects you obtain from the parser are lightweight containers for all the dense and ragged features that are part of an example graph, as well as the adjacency information. These can contain a single example graph or a batch of multiple graphs. You can access the tensors with an interface similar to that of Python dicts. For example, to access the age feature in the example, you would do this:

```
graph.node_sets["users"]["age"]

<tf.Tensor: shape=(4,), dtype=int32, numpy=array([24, 32, 27, 38], dtype=int32)>
```

Edge indices are accessed with their "adjacency" property:

```
graph.edge_sets["purchased"].adjacency.source

<tf.Tensor: shape=(7,), dtype=int32, numpy=array([0, 1, 2, 3, 4, 5, 5], dtype=int32)>
```

And similarly for context features:

```
graph.context["scores"]

<tf.Tensor: shape=(1, 4), dtype=float32, numpy=array([[0.45, 0.98, 0.1 , 0.25]], dtype=float32)>
```

A.2.2 Creating GraphTensors. Instances of `GraphTensor` can also be created as constants or from existing tensors. This is useful for writing unit tests and working in a Colab. To create a `GraphTensor`, you have to provide the various pieces and features forming the `GraphTensor`. This is best described by an example:

```
graph = tfgnn.GraphTensor.from_pieces(
    context=tfgnn.Context.from_fields(
        features={
            "scores": [[0.45, 0.98, 0.10, 0.25]],
        }),
    node_sets={
        "items": tfgnn.NodeSet.from_fields(
            sizes=[6],
            features={
                "category": ["food", "show ticket", "shoes",
                             "book", "flight", "groceries"],
                "price": tf.ragged.constant([[22.34, 23.42, 12.99],
                                             [27.99, 34.50],
                                             [89.99],
                                             [24.99, 45.00],
                                             [350.00],
                                             [45.13, 79.80, 12.35]]),
            }),
        "users": tfgnn.NodeSet.from_fields(
            sizes=[4],
            features={
                "name": ["Shawn", "Jeorg", "Yumiko", "Sophie"],
                "age": [24, 32, 27, 38],
                "country": ["usa", "uk", "japan", "france"],
            }),
    },
    edge_sets={
        "purchased": tfgnn.EdgeSet.from_fields(
            sizes=[7],
            features={},
            adjacency=tfgnn.Adjacency.from_indices(
                source=("items", [0, 1, 2, 3, 4, 5, 5]),
                target=("users", [1, 1, 0, 0, 2, 3, 0]),
            )),
        "is-friend": tfgnn.EdgeSet.from_fields(
            sizes=[3],
            features={},
            adjacency=tfgnn.Adjacency.from_indices(
                source=("users", [1, 2, 3]),
                target=("users", [0, 0, 0]),
            )),
    })
```

The data types for `tfgnn.Context`, `tfgnn.NodeSet`, `tfgnn.EdgeSet`, and `tfgnn.Adjacency` are exposed to the API for this construction to be possible. The classes verify that the ranks and shapes of the tensors are matching each other as constrained by the graph structure (i.e., in the same set they share a common prefix dimensions).

A.2.3 Reading GraphTensors. Graphs can be written to disk by encoding them into streams of `tf.train.Example` protos.

Graphs can be read into a TensorFlow program by decoding these `Example` protos to `GraphTensor` objects. To configure the parsing routine, the `GraphSchema` message is read and converted to a `GraphTensorSpec` object which describes the layout of the graph tensor in the TensorFlow runtime: the list of its various node sets, edge sets, and all the features associated with them and the context features. This data structure is analogous to the `tf.TensorSpec` object of TensorFlow. A `GraphTensorSpec` object is attached to all instances of `GraphTensor` and flows along with it. To read a stream of graph tensors, first create a spec, like this:

```
import tensorflow_gnn as tfgnn

graph_schema = tfgnn.read_schema(schema_path)
graph_spec = tfgnn.create_graph_spec_from_schema_pb(graph_schema)
```

You can then use the library's own parser to decode the `tf.train.Example` features into `GraphTensor` objects:

```
data_paths = tf.data.Dataset.list_files(...)
ds = tf.data.TFRecordDataset(data_paths)
ds = ds.batch(batch_size)
parser_fn = functools.partial(tfgnn.parse_example, graph_spec)
```

```
ds = ds.map(parser_fn)
```

The dataset ds yields a stream of instances of GraphTensor. The parsing function ingests a batch of serialized graphs (a rank-1 tensor of strings).

Similarly to tf.io.parse_single_example, there is also a function tfgnn.parse_single_example to parse an unbatched stream of encoded tf.train.Example strings with GraphTensor instance in them, and you can also run batch on top of that (i.e., you can batch GraphTensor instances):

```
data_paths = tf.data.Dataset.list_files(...)
ds = tf.data.TFRecordDataset(data_paths)
parser_fn = functools.partial(tfgnn.parse_single_example, graph_spec)
ds = ds.map(parser_fn)
ds = ds.batch(batch_size)
```

Typically this step is followed by a feature engineering step that normalizes, embeds and concatenates the various features into a single embedding for each set, and then batches multiple subgraphs into modified batched GraphTensor instances.

The containers can then be used to pick out various tensors and build a model, or simply to inspect their contents:

```
for graph in ds.take(10):
  tensor = graph.node_sets["item"]["category"]
  print(tensor)
```

You will go through important details of batching and flattening to a single graph below.

A.3 Broadcast and pool operations

Consider how the message passing API could help us to find the total user spending on purchased items. As a preparation step, let's calculate for each "item" the latest price value and materialize the result to the graph tensor.

```
item_features = graph.node_sets["items"].get_features_dict()
item_features["latest_price"] = item_features["price"][:, :1].values
graph = graph.replace_features(node_sets={"items": item_features})
print(graph.node_sets["items"])
```

User spending can now be calculated by first computing purchase prices by broadcasting item latest prices to the "purchase" edges. The total user spendings are computed by sum-pooling purchase prices to their users.

```
purchase_prices = tfgnn.broadcast_node_to_edges(
    graph, "purchased", tfgnn.SOURCE, feature_name="latest_price")
total_user_spendings = tfgnn.pool_edges_to_node(
    graph,
    "purchased",
    tfgnn.TARGET,
    reduce_type="sum",
    feature_value=purchase_prices)
print(total_user_spendings)
```

Note that API allows to reference feature values either by their name and by their values. For latter it is not required that feature exists in the graph tensor as soon as its shape is correct.

Message passing is also supported between the graph context and any node set or edge set. As an example, let's compare individual user spendings to the maximum amount spent by any users. The code below first max-pool all individual user spendings and then broadcast them back to "users" to compute fractions.

```
max_amount_spend = tfgnn.pool_nodes_to_context(
    graph, "users", feature_value=total_user_spendings,
    reduce_type="max")
max_amount_spend = tfgnn.broadcast_context_to_nodes(
    graph, "users", feature_value=max_amount_spend)

print(total_user_spendings / max_amount_spend)
```

A.4 Implementing GATv2

The following code snippet shows the essence of how GATv2 [7] has been implemented in TF-GNN. It performs attention from many senders to one receiver.

In the simplest case, the convolution is applied to an edge set whose target nodes are the receivers, and each receiver attends to the adjacent source nodes. However, TF-GNN lets you configure this is many ways: (1) The roles of source and target nodes can be swapped by setting the receiver_tag; for example, consider how an edge set of hyperlinks between HTML docs can reasonably used in either direction. (2) The sender value can be supplied by the neighbor node, the connecting edge, or both; this is controlled by the sender_*_feature args, at least one of which must not be None. (3) Attention can also happen with receiver_tag=tfgnn.CONTEXT and either all nodes or all edges in

the respective graph component as senders. This comes in handy for attention in the context update of a Graph Network [5] or for readout of a graph-level feature by smart pooling of node or edge states.

All those case distinctions are handled by the superclass, which passes the suitable broadcast and pool ops as arguments into the actual `GATv2Conv.convolve()` method.

```python
class GATv2Conv(tfgnn.keras.layers.AnyToAnyConvolutionBase):

  def __init__(self, *,
               num_heads, per_head_channels,
               receiver_tag=None,  # SOURCE, TARGET or CONTEXT.
               receiver_feature=tfgnn.HIDDEN_STATE,  # Required.
               sender_node_feature=tfgnn.HIDDEN_STATE,  # Set None to disable.
               sender_edge_feature=None,  # Set tfgnn.HIDDEN_STATE to enable.
               attention_activation="leaky_relu", activation="relu",
               edge_dropout=0., **kwargs):
    # Save initializer args.
    super().__init__(
        receiver_tag=receiver_tag, receiver_feature=receiver_feature,
        sender_node_feature=sender_node_feature,
        sender_edge_feature=sender_edge_feature,
        extra_receiver_ops={"softmax": tfgnn.softmax}, **kwargs)
    self._num_heads = num_heads
    self._per_head_channels = per_head_channels
    self._edge_dropout_layer = tf.keras.layers.Dropout(edge_dropout)
    self._attention_activation = tf.keras.activations.get(attention_activation)
    self._activation = tf.keras.activations.get(activation)
    # Create the transformations for the query input in all heads.
    self._w_query = tf.keras.layers.Dense(per_head_channels * num_heads)
    # Create the transformations for value input from sender nodes and edges.
    self._w_sender_node = self._w_sender_edge = None
    if self.takes_sender_node_input:
      self._w_sender_node = tf.keras.layers.Dense(per_head_channels * num_heads)
    if self.takes_sender_edge_input:
      self._w_sender_edge = tf.keras.layers.Dense(
          per_head_channels * num_heads,
          use_bias=self._w_sender_node is None)  # Avoid two biases.
    # Create the transformation to attention scores.
    self._attention_logits_fn = tf.keras.layers.experimental.EinsumDense(
        "...ik,ki->...i", output_shape=(num_heads,))

  def convolve(self, *,
               sender_node_input, sender_edge_input, receiver_input,
               broadcast_from_sender_node, broadcast_from_receiver,
               pool_to_receiver, extra_receiver_ops, training):
    # Form the attention query for each head.
    query = broadcast_from_receiver(self._split_heads(self._w_query(
        receiver_input)))
    # Form the attention value for each head.
    value_terms = []
    if sender_node_input is not None:
      value_terms.append(broadcast_from_sender_node(
          self._split_heads(self._w_sender_node(sender_node_input))))
    if sender_edge_input is not None:
      value_terms.append(
          self._split_heads(self._w_sender_edge(sender_edge_input)))
    value = tf.add_n(value_terms)
    # Compute the attention coefficients.
    attention_features = self._attention_activation(query + value)
    logits = tf.expand_dims(self._attention_logits_fn(attention_features), -1)
    attention_coefficients = extra_receiver_ops["softmax"](logits)
    attention_coefficients = self._edge_dropout_layer(attention_coefficients)
    # Apply the attention coefficients to the transformed query.
    messages = value * attention_coefficients
    pooled_messages = pool_to_receiver(messages, reduce_type="sum")
    # Apply the nonlinearity.
    return self._activation(self._merge_heads(pooled_messages))

  # The following helpers map forth and back between tensors with...
  #  - a separate heads dimension: shape [..., num_heads, channels_per_head],
  #  - all heads concatenated:     shape [..., num_heads * channels_per_head].
  def _split_heads(self, tensor):
    extra_dims = tensor.shape[1:-1]  # Possibly empty.
    new_shape = (-1, *extra_dims, self._num_heads, self._per_head_channels)
    return tf.reshape(tensor, new_shape)

  def _merge_heads(self, tensor):
    extra_dims = tensor.shape[1 : -2]  # Possibly empty.
    merged_dims = tensor.shape[-2:]
    new_shape = (-1, *extra_dims, merged_dims.num_elements())
    return tf.reshape(tensor, new_shape)
```

A.5 Orchestration

The following code snippet demonstrates a typical use of the *Orchestrator*. Quick start users need specify only:

(1) Training—and optional validation—dataset(s),
(2) Optional feature processing callbacks,
(3) A base Graph Neural Network. (A `tf.keras.Model` that both accepts and returns a GraphTensor),
(4) A `runner.Task`.

The snippet assumes a toy GraphSchema. The provided training—and optional validation—dataset(s) must contain either (i) serialized `tf.train.Example` protocol buffers to be parse by this schema or (ii) GraphTensor objects compatible with this schema. (As checked by `tfgnn.check_compatible_with_schema_pb`.) The toy GraphSchema:

```
context {
  features {
    key: "ni"
    value: {
      dtype: DT_INT64
      shape { dim { size: 1 } }
    }
  }
}
node_sets {
  key: "coconut"
  value {
  }
}
node_sets {
  key: "grail"
  value {
    features {
      key: "holiness"
      value {
        dtype: DT_INT64
        shape { dim { size: 1 } }
      }
    }
  }
}
node_sets {
  key: "swallow"
  value {
    features {
      key: "speed"
      value {
        dtype: DT_FLOAT
        shape { dim { size: 1 } }
      }
    }
  }
}
edge_sets {
  ...
}
edge_sets {
  ...
}
edge_sets {
  ...
}
edge_sets {
  ...
}
```

And *orchestration*:

```python
import tensorflow as tf
import tensorflow_gnn as tfgnn

from tensorflow_gnn import runner

gtschema = tfgnn.read_schema(".../schema.pbtxt")
gtspec = tfgnn.create_graph_spec_from_schema_pb(gtschema)

train_ds_provider = runner.TFRecordDatasetProvider(".../train*")

# Extract labels from the graph context, importantly: this lambda
# matches the `runner.GraphTensorProcessorFn` protocol.
def extract_label_fn(gt):
    return gt, gt.context["ni"]

# Process node features, importantly: this callback is executed as
# a part of a `tf.data.Dataset`. (Execution takes place on CPU and
# not any accelerator.)
def node_sets_processor_fn(node_set, *, node_set_name):
  if node_set_name == "coconut":
    return {"empty": tfgnn.keras.layers.MakeEmptyFeature(node_set)}
  if node_set_name == "grail":
    return {"holiness": tf.keras.layers.Hashing(153)(node_set["holiness"])}
  if node_set_name == "swallow":
```

```python
    return {"speed": tf.math.log1p(node_set["speed"])}

# Set initial node states, importantly: this callback is executed
# as a part of modeling. (Learnable weights are included in back
# propagation.)
def initial_node_states_fn(node_set, *, node_set_name)
  if node_set_name == "coconut":
    return node_set["empty"]
  if node_set_name == "grail":
    return tf.keras.layers.Embedding(153, 16)(node_set["holiness"])
  if node_set_name == "swallow":
    return tf.keras.layers.Dense(16)(node_set["speed"])

def model_fn(gtspec):
  # simple_gnn is a function: Callable[[tfgnn.GraphTensorSpec], tf.keras.Model]. Where
  # the returned model both takes and returns a scalar `GraphTensor`
  # for its inputs and outputs.
  return tf.keras.Sequential([
      tfgnn.keras.layers.MapFeatures(node_sets_fn=initial_node_states_fn),
      simple_gnn(gtspec),
  ])

# Binary classification by the root node.
task = runner.RootNodeBinaryClassification(node_set_name="grail")

trainer = runner.KerasTrainer(
    strategy=tf.distribute.TPUStrategy(),
    model_dir=".../monty/",
    # len(train_ds_provider.get_dataset(...)) == 8191 and global_batch_size == 128.
    steps_per_epoch=8191 // 128,
    restore_best_weights=False)

runner.run(
    train_ds_provider=train_ds_provider,
    train_padding=runner.FitOrSkipPadding(gtspec, train_ds_provider),
    model_fn=model_fn,
    optimizer_fn=tf.keras.optimizers.Adam,
    epochs=4,
    trainer=trainer,
    task=task,
    gtspec=gtspec,
    global_batch_size=128,
    feature_processors=[
        extract_label_fn,  # Extract any labels first.
        tfgnn.keras.layers.MapFeatures(node_sets_fn=node_sets_processor_fn),
    ])
```

A.6 Case Study: OGBN-MAG

Full definitions for files used in the case study on OGBN-MAGG[9] (Section 8) are provided here for completeness.

A.6.1 Graph Schema. This is the protocol buffer definition of the Graph Schema used for OGBN-MAG from Figure 5.

```
node_sets {
  key: "author"
  value {
    features {
      key: "#id"
      value {
        dtype: DT_STRING
      }
    }
    metadata {
      filename: "nodes-author.tfrecords@15"
      cardinality: 1134649
    }
  }
}
node_sets {
  key: "field_of_study"
  value {
    features {
      key: "#id"
      value {
        dtype: DT_STRING
      }
    }
    metadata {
      filename: "nodes-field_of_study.tfrecords@2"
      cardinality: 59965
    }
  }
}
node_sets {
```

---

[9]https://ogb.stanford.edu/docs/nodeprop/#ogbn-mag

```
  key: "institution"
  value {
    features {
      key: "#id"
      value {
        dtype: DT_STRING
      }
    }
    metadata {
      filename: "nodes-institution.tfrecords"
      cardinality: 8740
    }
  }
}
node_sets {
  key: "paper"
  value {
    features {
     key: "#id"
      value {
        dtype: DT_STRING
      }
    }
    features {
      key: "feat"
      value {
        dtype: DT_FLOAT
        shape {
          dim {
            size: 128
          }
        }
      }
    }
    features {
      key: "labels"
      value {
        dtype: DT_INT64
        shape {
          dim {
            size: 1
          }
        }
      }
    }
    features {
      key: "year"
      value {
        dtype: DT_INT64
        shape {
          dim {
            size: 1
          }
        }
      }
    }
    metadata {
      filename: "nodes-paper.tfrecords@397"
      cardinality: 736389
    }
  }
}
edge_sets {
  key: "affiliated_with"
  value {
    source: "author"
    target: "institution"
    metadata {
      filename: "edges-affiliated_with.tfrecords@30"
      cardinality: 1043998
    }
  }
}
edge_sets {
  key: "cites"
  value {
    source: "paper"
    target: "paper"
    metadata {
      filename: "edges-cites.tfrecords@120"
      cardinality: 5416271
    }
  }
}
edge_sets {
  key: "has_topic"
  value {
    source: "paper"
    target: "field_of_study"
```

```
      metadata {
        filename: "edges-has_topic.tfrecords@226"
        cardinality: 7505078
      }
    }
  }
}
edge_sets {
  key: "writes"
  value {
    source: "author"
    target: "paper"
    metadata {
      filename: "edges-writes.tfrecords@172"
      cardinality: 7145660
    }
  }
}
```

A.6.2 Generated Sample Spec. This is the protocol buffer generated for input to a sampling pipeline from the code snippet in Figure 6.

```
seed_op {
  op_name: "SEED->paper"
  node_set_name: "paper"
}
sampling_ops {
  op_name: "paper->paper"
  input_op_names: "SEED->paper"
  edge_set_name: "cites"
  sample_size: 32
  strategy: RANDOM_UNIFORM
}
sampling_ops {
  op_name: "(paper->paper|SEED->paper)->author"
  input_op_names: "paper->paper"
  input_op_names: "SEED->paper"
  edge_set_name: "written"
  sample_size: 8
  strategy: RANDOM_UNIFORM
}
sampling_ops {
  op_name: "author->institution"
  input_op_names: "(paper->paper|SEED->paper)->author"
  edge_set_name: "affiliated_with"
  sample_size: 16
  strategy: RANDOM_UNIFORM
}
sampling_ops {
  op_name: "author->paper"
  input_op_names: "(paper->paper|SEED->paper)->author"
  edge_set_name: "writes"
  sample_size: 16
  strategy: RANDOM_UNIFORM
}
sampling_ops {
  op_name: "(author->paper|SEED->paper|paper->paper)->field_of_study"
  input_op_names: "author->paper"
  input_op_names: "SEED->paper"
  input_op_names: "paper->paper"
  edge_set_name: "has_topic"
  sample_size: 16
  strategy: RANDOM_UNIFORM
}
```

A.6.3 Hyper-parameter optimization. In their Vizier study, our Googler sets up the following hyperparameter search problem:
 (1) 'message_dim': the dimension of hidden states for each node. Discrete Search space: '[32, 64, 128, 256, 512]'
 (2) 'reduce_type': the strategy to pool the messages from edges to receiver nodes. Categorical Search space: '["sum", "mean"],
 (3) 'l2_regularization': the coefficient of l2 regularization for weights and biases. Continuous (Log) Search Space: '[1e-6, 1e-4].
 (4) 'dropout': the dropout probability. Discrete Search space: '[0.1, 0.2, 0.3]'
 (5) 'use_layer_normalization': Flag to determine whether to utilize layer normalization when updating the node states. Boolean Search space: '[True, False]'

They use the Adam optimizer with a cosine decay learning rate schedule.

Results. The top-3 performing models on the validation set had 'message_dim = 256', 'reduce_type = "sum"', 'dropout = 0.2', 'use_layer_normalization = True' while the 'l2_regularization' varied in the range [1e-6, 4e-6].

A.6.4 A minimal usage of the TF-GNN Orchestration Layer. This illustrates how to train the model defined in Section 7 using the Orchestrator.

```
from tensorflow_gnn import runner

train_ds_provider = runner.TFRecordDatasetProvider("/home/...")
feature_processors = [extract_labels, feature_mapping]  # See appendix.
```

```python
task = runner.RootNodeMulticlassClassification(
    node_set_name="paper", num_classes=349)

trainer = runner.KerasTrainer(
    strategy=tf.distribute.MirroredStrategy(), model_dir="/tmp/my_model",
    steps_per_epoch=num_training_examples // global_batch_size,
    restore_best_weights=False)

optimizer_fn = functools.partial(
    tf.keras.optimizers.Adam,
    learning_rate=tf.keras.optimizers.schedules.CosineDecay(...))

runner.run(
    train_ds_provider=train_ds_provider,
    gtspec=tfgnn.create_graph_spec_from_schema_pb(graph_schema),
    feature_processors=feature_processors,
    task=task, model_fn=model_fn, # Model as defined above.
    trainer=trainer, optimizer_fn=optimizer_fn,
    epochs=20, global_batch_size=global_batch_size)
```